

Applying multiagent simulation to planetary surface operations

Maarten Sierhuis¹
Michael H. Sims²
William J. Clancey²
Pascal Lee³

1. *RIACS/USRA, NASA/Ames Research Center, MS 19-39, Moffett Field, CA 94035-1000*
2. *NASA/Ames Research Center, MS 269-3, Moffett Field, CA 94035-1000*
3. *SETI Institute, NASA/Ames Research Center, MS 245-3, Moffett Field, CA 94035-1000*

Abstract. This paper describes a multiagent modeling and simulation approach for designing cooperative systems. Issues addressed include the use of multiagent modeling and simulation for the design of human and robotic operations, as a theory for human/robot cooperation on planetary surface missions. We describe a design process for cooperative systems centered around the Brahms modeling and simulation environment being developed at NASA Ames.

Keywords. Human-centered design, human-activity models, multiagent models, multiagent simulation.

1. Design problem and objectives

The establishment of remote field camps is a proven strategy in geographic exploration on Earth and is likely to be a required capability on human missions to Mars to extend the range and safety of field exploration activities. Robotics could play a key role in helping support this need. From the start it seems clear that robots will cooperate with humans [1]. The question is, to what extent will human-robot cooperation be necessary, and how will this take place? We are in the process of starting a detailed study of scenarios for human/robotic cooperation in establishing remote field camps for excursions on Mars.

Cooperative work practice scenarios for establishing and using of remote field camps are being developed. These scenarios will be computationally investigated in the Brahms modeling and simulation environment. In addition, we will do field experiments to verify appropriateness and utility of the Brahms models as a tool for planning and designing human/robotic cooperative activities on Mars. This verification will include evaluating that activity times are realistic, that resource use is consistent with the model's assumed constraints and the assumptions of communication activities are enabling of the proposed field camp requirements. We anticipate the specifics of this research to be interesting to the Mars Mission planning community, however our fundamental objective is to develop tools for designing cooperative autonomous systems and to show the utility of those tools and assumptions (such as the needed levels of autonomy) in a high fidelity, realistic evaluation of robotic activities on Mars.

The objectives of this research are:

- To establish appropriate robot and human allocation of activities in the establishment of habitat structures on a planetary surface.
- To establish the utility of the Brahms environment for mission planning by demonstration on a Mars-relevant scenario—the establishment of a remote science field site.
- To establish guidelines for judicious use of robots leading to human risk reduction.

- To establish a framework for management and optimization of robotic colonies for planetary surface tasks associated with human presence.

2. State of the art

Mission planning for robotic and mixed human-robotic tasks is currently done quite informally with the design team's heuristic intuitions about tasks the agents (either human or robotic) need to do and the likelihood of that capability being available in the future state of the art. This creates a fundamental problem with the analysis of the robotic elements of a mission being carried out at a very high level of abstraction until well into the commitment for a mission. In part this is a consequence of the inadequacy of current systems in allowing easy modeling of the intricacies of a rich and dynamic set of activities being carried out by robots in conjunction with humans. This work is focused on directly alleviating this problem.

It is useful to study a dynamic real-world system to learn something about its behavior. However, often it is necessary to use a model of the system to study its performance, since experimentation with the system itself would be disruptive, not cost-effective or simply impossible because the system hasn't been developed yet. In manufacturing design, computational simulation tools have been in use for several years [2]. However, the state-of-the-art for such tools is to use discrete-event simulation generally stochastic in nature. In such models, variables that represent the "soft" nature of the system—such as arrival times of jobs at a point in the system—can only be modeled stochastically. Often, it is these types of variables that are trying to capture how things really work in the real world. Especially in soft systems [3], i.e. systems where human activity, communication and cooperation—the work practice—play an important role in the performance of the system, it is very difficult to develop a good model using probabilistic behavior. We address this problem by using a *qualitative symbolic model-based simulation approach*.

The design problem we are addressing, i.e. the amount of human-robot cooperation needed in the system, has all the elements of a soft system. The use of modeling and simulation to the design and understanding of the cooperative system is the right approach, because there is no current real-world system to experiment with. Also, the state-of-the-art robotic systems are not yet capable of the kinds of autonomy and human cooperation that will be needed for this kind of task. This makes it difficult to test our design in field experiments without the guidance of detailed models.

3. The epistemological level of work practice

We briefly describe our theory of modeling work practice. Representing how people do work can be done at many different levels. In the knowledge engineering and AI-world, people's work has been described in terms of their problem-solving expertise. The theory is that we can model people's problem-solving behavior by representing this behavior in a computational qualitative model that is able to duplicate some of this behavior. Work process models, such as Petri-Net models of a work process, describe what tasks are performed and when—i.e. transitions. In workflow models we describe how a specific product "flows" through an organization's work process. This describes the sequential tasks in the work process that "touch" a work-product. All these modeling approaches describe the work in an organization at a certain level of detail. However, what is missing from all these modeling approaches is a representation of how work gets done. What is missing is a description of the work at the *work practice level*.

Work practice relates situatedness and how things work in real-life to abstract rules in methods and procedures. To understand how the performance of procedures in practice differs from the abstract methods and procedures specified in designs we model the cooperative activities of agents as they occur within the context of the actual work practice.

Work practice includes those aspects of the work that make people behave a certain way in a specific situation, and at a specific moment in time. To describe people's situation-specific behavior we need to include those aspects of the situation that explain the influence on the *activity behavior* of individuals (in contrast with problem-solving behavior). Following is a brief description of some important aspects that determine an individual's situation-specific behavior.

Activity behavior

People's behaviors are emergent from the "execution" of specific activities at certain moments. A person or system cannot be "alive" without being in some kind of activity. Even "doing nothing" is described in terms of a "do-nothing" or idle activity. Furthermore, what activity is being performed depends on the situational context that a person or system is in. Agents' behaviors are organized into activities, inherited from groups to which agents belong. Most importantly, activities locate behaviors of people and their tools in time and space, so that resource availability and informal human participation can be taken into account [4].

Activities can be subsumed by other activities in a hierarchical structure. With this we mean that a person can be in multiple subsumed activities at once. For example, you can be in the activity of reading a book, while at the same time be in the higher level activity of a being on a business trip. When the phone rings in your hotel room, you get up and walk over to pick up the phone. This means that you interrupt the activity of reading your book, and start the activity of answering the phone. You actually never stop being in the activity of reading your book, but you merely suspend the activity to focus on a new activity, continuing with the suspended activity when the phone call is over.

A model of activities does not necessarily describe the intricate details of reasoning or calculation, but instead captures aspects of the social-physical context, including space and time in which reasoning occurs [5] [6]. Activities subsume goals, and goal-directed behavior occurs within an activity. Therefore, an activity-based model does not necessarily include the goals of the agent. It is a model at the social-level, relating the collaboration and communication of agents with the interaction of these agents in the real world.

Context

People act based on the situation they are in [7]. With this we mean that people do not have a rigid pre-specified plan that they are following, but that they behave based on their beliefs about what they experience (infer or detect) their context to be. Different people can/will have different beliefs about a similar context. If we want to model work practice, we need to be able to separate the context from people's different interpretation of that context. In order to do so, we describe context in terms of objects and artifacts that people observe and use within their environment. We also describe the geographical locations of people and artifacts. What describes a context is known as world-facts or simply facts. Facts represent factual information about the three-dimensional world people live in. People do not automatically have "knowledge" about those facts, and if people have "knowledge" about those facts it might not be correct. For example, you can believe that your car is parked in the garage, whereas in reality someone has taken the car to go out. So, the fact is that the location of your car is wherever it has been taken, while you believe that the location of the car is the garage. You will have that belief until either someone tells you about the actual location (or wrong location) of the car, or until you go to the garage and observe (i.e. detect) that the car is not there. Of course, if the car is returned before any of this takes place you will never know the car had been gone. In other words, although facts are global (the car can only be in one location), not every person can get "access" (i.e. get a belief) about that fact.

Implicit in the above example is the fact that people and objects are always located and moving from one location to another.

Communication

An important aspect of cooperation is that people communicate with each other. Theoretically, we can define communication as the transfer of beliefs from one person or object to another [8]. However, there are many types of communication that people use: face to face, phone, e-mail, fax, voice loop, et cetera. Each of these types of communication has a different practice. For example, in communication via e-mail the sender and receiver do not have to be in the same geographical location and time zone. The use of specific tools and artifacts are an important enabler or constraint on the effectiveness and efficiency of the communication, thus impacting how we work and collaborate. Therefore, to understand work practice, we need to understand not only when and what people communicate, but also what type of communication is used and how artifacts are used in the communication activity.

Communities of practice

In order to describe how two different persons can perform different activities based on the same situational context, we borrow the term community of practice (CoP) from the social sciences [9]. People belong to many different communities. One way we can distinguish one community from another is in the way they are able to perform certain activities. For instance, at NASA we can distinguish the community of Apollo astronauts from the rest of the communities at NASA. We can describe the work of a particular community as a separate "group." Members of groups can perform the group's activities. Thus, we can describe people's behavior in terms of the groups they belong to.

4. Brahms: activity-based multiagent modeling

A traditional task or functional analysis of work leaves out the logistics, especially how conditions come to be detected and resolved, such that work and information actually flows. Without this understanding we cannot properly design intelligent agents that automate human tasks or interact with people as their collaborators. What is wanted is a model that includes aspects of reasoning found in an information-processing model, plus aspects of geography, agent movement, and physical changes to the environment found in a multiagent simulation [10]. A model of work practice focuses on informal, circumstantial, and located behaviors by which *synchronization* occurs, such that the task contributions of humans and machines flow together to accomplish goals.

Brahms is a multiagent simulation program developed by ([1]) that allows the explicit modeling of activities of people and systems. The approach is qualitative, and relates knowledge-based models of cognition (e.g., task models) [12] with discrete-event simulation and a behavior-based subsumption architecture [13] [14] [5] [15].

Agents' behaviors are organized into activities, inherited from groups to which agents belong. Groups include not only technical functions (such as "Shuttle tile specialist"), but also where people work ("Orbiter Processing Facility people"), their temporary roles ("Atlantis flight prep coordinator"), their background ("USA contractor, previously at Boeing"), and the tools they use ("XYZ database users"). Most importantly, activities locate behaviors of people and their tools in time and space, such that resource availability and informal human participation can be taken into account.

Thus Brahms differs from other multiagent systems by incorporating the following:

- *Chronological activities of multiple agents* — attention and cooperation is modeled according to simultaneous participation in different groups (identities), determining what

is perceived and how it is interpreted; behaviors are chunked according to how agents allocate time during the day and use different spaces.

- *Conversations* at the level of sequences of ask/tell interactions, reading and writing documents and databases (e.g., using speech to control robots during an extra-vehicular activity).
- *How information is represented, transformed, reinterpreted in various physical modalities* — online manuals, databases, forms, multiple pass reviewing and reading, location and movement of documents (e.g., procedures on clip boards in the Station).
- *Multiple graphic views of work* (e.g., geographic layout, agent-centric [chronological schedules & checklists], job-centric [workflow diagrams]) amenable for use by engineers, planners, scientists, and managers, especially across organizations (e.g., in payload processing, relating the university PI to MSFC designers to JSC trainers).

A Brahms model can be used to simulate human-machine systems for what-if experiments, for training, for “user models,” or for driving intelligent assistants and robots. The architecture includes the following (simplified) representational constructs:

```

Groups of groups containing
  Agents who are located and have
    Beliefs that lead them to engage in
      Activities that are specified by
        Workframes that consist of
          Preconditions of beliefs that lead to
            Actions, consisting of
              Communication Actions
              Movement actions
              Primitive Actions
              Other composite activities
            Consequences of new beliefs and world facts
          Thoughtframes that consist of
            Preconditions and
            Consequences
  
```

In addition, *active physical objects* (e.g., cameras, telephones, laptop computers are modeled as entities whose state can also change by the application of workframes and thoughtframes. *Conceptual objects* are entities people have beliefs about, but that have no specific location (e.g., a mission) and are associated with physical objects (e.g., a particular orbiter)

5. Methodological approach

The cooperative design approach we are proposing is represented in Figure 1. This shows a flowchart of the type of output and the processes. This design, simulate and test approach allows us to make a number of cycles, and improve on designs. Although our methods may be used for optimization, the goal is to find single point solutions for which we can convincingly demonstrate those solutions as consistent with Mars mission constraints. For Example, in our research on remote field camps on Mars the technical activities will be:

1. Developing a deep understanding of the problem of establishing and a use of a Mars remote science outpost.
2. Developing models of the robotic and human activities associated with that remote science outpost (using the Brahms tool as a design capture tool for cooperation).
3. Simulate these models computationally leading to activity timelines and communication and other constraint consistencies.
4. Cycle back to 1 and 2 as necessary.

5. Given realistic scenarios, actors, actions and communications as output to the above process, we will next do field experiments to convincingly demonstrate the realism of these.
6. Cycle back to 1 and 2 as necessary.

These activities are represented graphically in Figure 1.

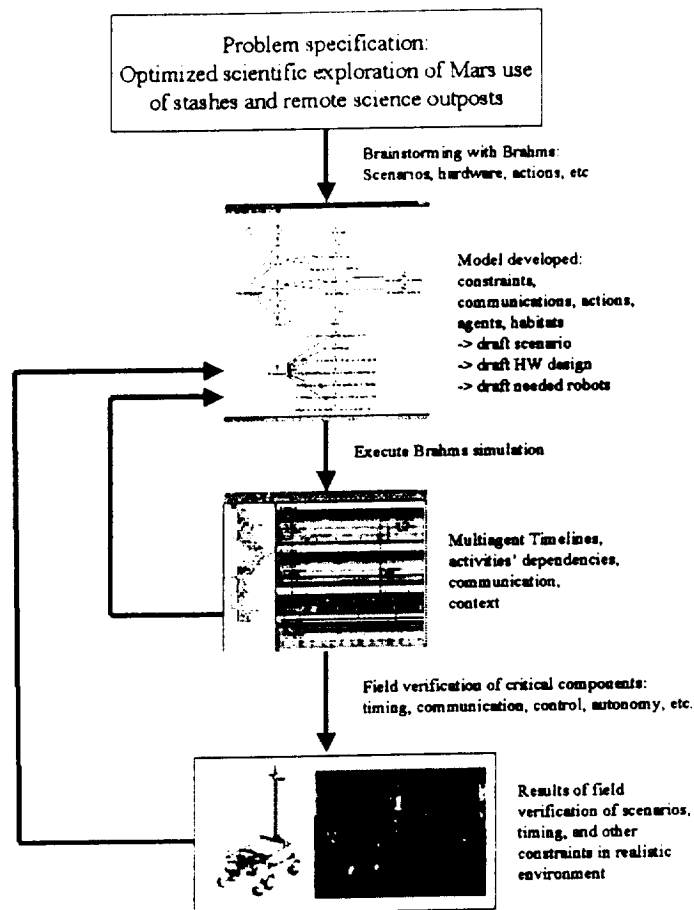


Figure 1. Work activity flowchart

6. Simulation process

Many have described the process of a successful simulation [2] [16] [17]. All of them mention a series of processes that need to be followed. The high-level processes are shown in Figure 2.




Figure 2. Simulation modeling process (borrowed from [17])

A simulation study first starts with understanding the *real world*, as well as the problem to be tackled. In our study, the "real world" is that of robotic deployment of stashes and remote science outposts on Mars. The problem to be tackled is that of the *design* of such a mission using cooperative autonomous robots, for the *purpose* of understanding the limitations of autonomous robots and the extent we need human-robot cooperation to accomplish the tasks. To design the "real world" we start with a *conceptual modeling* activity. For this study, we use a qualitative static modeling approach called World Modeling (WM) [18]. The output of this process is a detailed model of activities for the tasks involved, the distribution of those activities over humans and robots, as well as a set of constraints for these activities (environmental, communication, timing, cooperation, habitat, and robot design).

Figure 3 shows a part of a WM conceptual model from a simulation project in which we simulated the cooperative activities of the Apollo 12 astronauts in deploying the Apollo Lunar Surface Experiment Package (ALSEP) [19]. It shows the *Remove ALSEP Package-1 Activity* the Commander (CDR) Pete Conrad performs. It includes the timing of this activity, decomposed models of all the sub-activities (see the right side of Figure 3), voice-loop communications, geographical location, and tools to be used.

After this, the model has to be coded into a *computer model*. This is done in the Brahms environment. When the model is complete, experiments are run to develop *solutions* to the real world problem being handled. Doing this, a design solution of all the robot/human activities for the task at hand is obtained. In Brahms, solutions come in the form of a high-fidelity simulation of all the cooperative agents, artifacts and environmental constraints relevant for the activities during the mission. The end-user can analyze the simulation output in the form of a multiagent activity timeline (2D display), and a database of historical simulation data that can be used for statistical analysis.

Figure 4 shows the multiagent activity timeline from the simulation of the conceptual model of the Apollo12 ALSEP Deployment from Figure 3. This figure shows the two lunar surface astronauts, Lunar Module Pilot (LMP) Al Bean (top), CDR Pete Conrad (2nd from the top) and the Capsule Communicator (CapCom) Ed Gibson, located at the Manned Spaceflight Center (bottom). The arrows show the communication over the voice loop, including the simulation of time delay to/from Earth (agent 2nd from the bottom).

The solutions found in the simulation experiments can be implemented in the real world, and a better understanding of the problem will lead to better decision making in the design of collaborative robots. We will implement the model in robotic field experiments. These

experiments will allow us to validate the models, allowing us to validate the cooperation and communication between humans and robots, which in turn allows us to enhance the models in subsequent phases of the project.

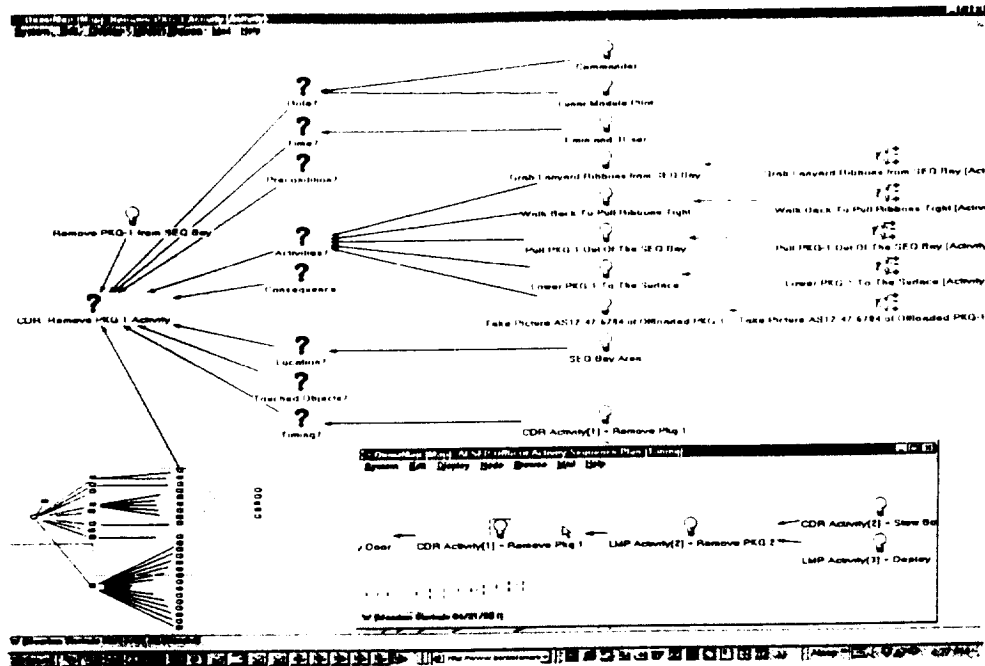


Figure 3. Conceptual Model of Apollo 12 ALSEP Deployment Activity

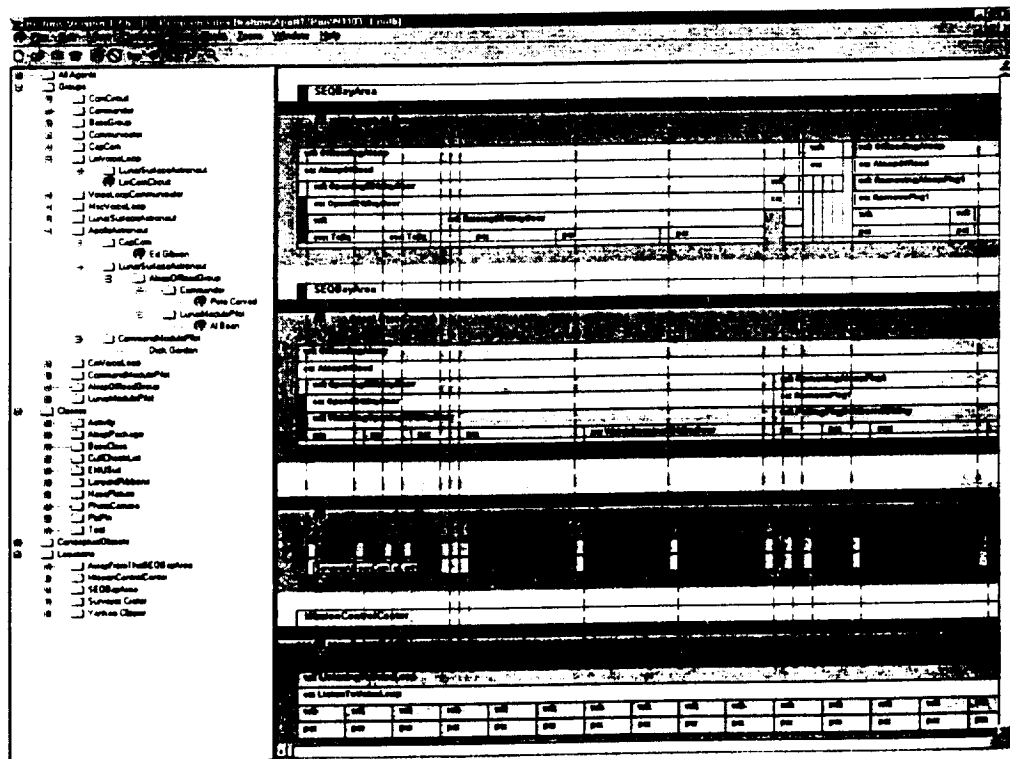


Figure 4. Multiagent Activity Timeline Output of the Brahms Simulation of the Apollo 12 ALSEP Deployment

7. Human errors

The notion of human error is often cited as the cause of mistakes and disasters. What constitutes a human error is often left to the interpretation of the accident investigator. Usually, a deviation from a standard or nominal procedure is identified as an "error." However, when we consider the situation specific issues and the work practice in contrast to the procedures, we often find that the human error lies in the design and validation process of the procedures, i.e. the human errors are created by the engineers, not the crew. Very often it are the procedures that do not take into account the context and the situation in which the activities take place. Were the right people involved in the planning process? How were interactions between subsystems tested (consider the flaw of the recent Mars Polar Lander)? Did organizational/functional breakdown prevent interactions between activities, materials and the situation from being properly modeled?

The following "human error" occurred during the Apollo 16 deployment of the Heath-Flow Experiment (HFE) [19]. The HFE deployment was performed as part of the ALSEP deployment, the main task during the first EVA. The LMP was in the process of drilling a hole in the lunar surface to implant the first HFE probe. He had connected the HFE package to the Central Station (C/S) with a flatbed cable. At the same time the CDR was busy deploying the Passive Seismic Experiment (PSE) in close proximity of the C/S. All this was planned and trained and had very detailed procedures. Unfortunately, although known at the time, the procedures and training did not include the fact that the flatbed cables would not lay flat on the lunar surface due to the minimal lunar gravity. For example, the procedures did not include specific instructions on how to avoid getting tangled in one of the cables—it was very difficult for the astronaut to see his feet through the visor of his helmet. Consequently, the cable connecting the HFE to the C/S got hooked on one of the CDR's boots without the CDR noticing it, thus ripping the cable of the C/S and breaking the connection, making the Apollo 16 HFE unusable.

However, if we consider the CDR's specific situation, the procedures and his training displaying itself through the work practice of the astronauts, it becomes obvious that we should not call this an "astronaut error." The CDR's actions were not a deviation of the nominal procedures, nor an unintentional mistake. It was the situation specific context on the moon that showed the error in the procedures and designs, as well as the lack of work practice on the moon. Procedure designers and HFE engineers did not take this into account.

Similar problems will undoubtedly manifest themselves in future cooperation between humans and humans and autonomous robots. Nominal procedures will not capture the intricacies of the human-robot work practice. One of the benefits of modeling and simulating not just the nominal procedures, but also the work practice of how the procedures are put into action, including the effects of the environment, communication, tools and artifacts, and error conditions is that we can be more detailed in the design of the activities and interaction of the agents with each other and the environment. Thus, avoiding the lack of contextual (nominal) procedures and increasing the descriptions of how activities will be performed in reality, therefore lowering the chance of unplanned activities causing problems. Note that without considering these issues, Brahms models could incorporate the same kinds of failures. Therefore, a solid engineering framework is required, by which we can include systematical failure analysis of past designs (of which the HFE is one example). In our research at NASA we are working create a human-activity modeling and simulation methodology that can root out these problems in advance.

8. Conclusions

In this paper we described a multiagent modeling and simulation approach to the design of human-activity systems in general, and cooperative human-robotic activity systems for Mars missions in particular. The described approach is methodologically speaking a model-based approach, in combination with a more pragmatic approach in which field tests are performed to seek feedback on model validity. We propose this as a *human-centered design* approach that allows designers to include the aspects of work practice into the design of cooperative systems.

At the center of modeling the cooperative system is the Brahms activity-based multiagent modeling and simulation environment. The Brahms modeling language was specifically designed to describe the work practice of people and systems in the situated environment. Brahms models the situated activity behavior of each individual agent within its environment, allowing to model situated action. As part of our research we are at the start of applying this approach to the design of a cooperative system of humans and robots deploying remote scientific field camps on Mars.

9. References

- [1] P. D. Spudis, "Robots vs. Humans: Who Should Explore Space?," in *Scientific American*, vol. Vol. 10, 1999, pp. 24-31.
- [2] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis (2nd Edition)*. New York: McGraw-Hill, 1991.
- [3] P. Checkland and J. Scholes, *Soft Systems Methodology in Action*. Chichester, England: John Wiley & Sons Ltd., 1990.
- [4] B. A. Nardi, "Context and Consciousness: Activity Theory and Human-Computer Interaction," . Cambridge, MA: The MIT Press, 1996.
- [5] W. J. Clancey, "The Conceptual Nature of Knowledge, Situations, and Activity," in *Human and Machine Expertise in Context*, P. Feltyovich, R. Hoffman, and K. Ford, Eds. Menlo Park, CA: The AAAI Press, 1997, pp. 247-291.
- [6] J. Greenbaum and M. Kyng, "Design at Work: Cooperative design of computer systems," . Hillsdale, NJ: Lawrence Erlbaum, 1991.
- [7] L. A. Suchman, *Plans and Situated Action: The Problem of Human Machine Communication*. Cambridge, MA: Cambridge University Press, 1987.
- [8] J. Searle, R., *Speech Acts*. Cambridge, UK: Cambridge University Press, 1969.
- [9] E. Wenger, *Communities of Practice; Learning, meaning, and identity*: Cambridge University Press, 1997.
- [10] M. Tokoro, "Proceedings of the Second International Conference on Multi-Agent Systems," presented at Proceedings of the Second International Conference on Multi-Agent Systems, Kyoto, Japan, 1996.
- [11] W. J. Clancey, P. Sachs, M. Sierhuis, and R. van Hoof, "Brahms: Simulating practice for work systems design," *International Journal on Human-Computer Studies*, vol. 49, pp. 831-865, 1998.
- [12] J. E. Laird, A. Newell, and P. S. Rosenbloom, "Soar: An architecture for general intelligence," *Artificial Intelligence*, vol. 33, pp. 1-64, 1987.
- [13] R. Brooks, A., "Intelligence without representation," *Artificial Intelligence*, vol. 47, pp. 139-159, 1991.
- [14] W. Clancey, J., *Situated Cognition: On Human Knowledge and Computer Representations*: Cambridge University Press, 1997.
- [15] H. Nakashima, L. Noda, and K. Handa, "Organic Programming language GAEA for multi-agents," presented at Proceedings of the Second International Conference on Multi-Agent Systems, Kyoto, Japan, 1996.
- [16] J. Banks, J.S. Carson, and B. L. Nelson, *Discrete-Event System Simulation, 2nd Edition*. Upper Saddle River, NJ: Prentice-Hall, 1996.
- [17] S. Robinson, "Simulation Verification, Validation and Confidence: A Tutorial," *Transactions of The Society for Computer Simulation International*, vol. Vol. 16, pp. 63-69, 1999.
- [18] M. Sierhuis and A. M. Selvin, "Towards a framework for collaborative modeling and simulation," presented at presented at the Workshop on Strategies for Collaborative Modeling and Simulation, CSCW '96, Boston, MA, 1996.
- [19] E. Jones, M., "The Apollo Lunar Surface Journal," National Aeronautics and Space Administration, WWW URL: <http://www.hq.nasa.gov/office/pao/History/alsj/>, 1996 1997.